

- 1 -

METHOD AND DATA PROCESSING SYSTEM WITH DATA REPLICATION

CROSS-REFERENCE TO RELATED APPLICATION

This application relates to and claims priority from Japanese Patent Application No. 2003-108361, filed on April 11, 2003, the entire disclosure 5 of which is incorporated herein by reference.

BACKGROUND OF THE INVENTION

The present invention relates to a database processing technique used in operation of a database.

The online business performing the conventional large amount of transaction processing requires daily, weekly, and monthly processing such as a large amount data calculation and batch registration or a periodical batch processing such as database backup and re-organization. Since these processes are accompanied by batch access of a database used by the online service, the online business is greatly affected by it and it disturbs provision of 24-hour continuous service (function realized by execution of a program and an object).
10
15

Recently, the SAN (Storage Area Network) configuration organically connecting a plurality of storage devices such as magnetic disc devices by a dedicated high-speed network is in widespread use. In this configuration, the storage device provides a
20

function of high-speed copying of an arbitrary logical volume to a plurality of logical volumes, a function of data multiple writing by making an arbitrary logical volume a main volume and the other plurality of logical
5 volumes sub-volumes, a function of disconnecting the logical volume in the multiple write mode at an arbitrary moment and accessing the main volume and the sub-volumes independently of each other, and a difference reflection function for making an arbitrary
10 volume a main volume and another logical volume a sub-volume and resuming the multiple write so that only a difference between the data is reflected from the main volume to the sub-volume.

Moreover, a duplicate database access device
15 has been suggested in US Patent Application Publication number US 2002/0198899A1. When one of the database areas multiple-written or duplicated by the external storage device is disconnected and a physical duplicate database is created, it is possible to maintain
20 consistency of the duplicate database while continuing the transaction processing and to access a plurality of physical databases without modifying the application program which has been accessing the duplicate origin database.

25 SUMMARY OF THE INVENTION

In the aforementioned conventional technique, a duplicate of a database is created as a duplicate

database. When the database should be subjected to a predetermined processing such as reorganization and fragmentation, the object to be operated is switched to the duplicate database. When database processing is continued by the duplicate database even when reorganizing the database, update processing performed to the duplicate database should also be performed to the database after reorganization. However, in order to realize it, the update processing for the duplicate database should be reflected in time series to the database. Update processes performed by a plurality of programs such as a business program and an application program cannot be easily re-executed in time series by the aforementioned program. This has been difficult to be realized by other than the database managing side.

It is therefore an object of the present invention to solve the aforementioned problem and provide a technique capable of executing a database processing and a predetermined processing in parallel and concurrently when the predetermined processing should be performed to the database.

According to the present invention, an object to be operated is switched to a duplicate database so that database processing is continued and a predetermined processing is executed to a main database. After the update processing performed to the duplicate database is reflected on the main database, the object to be operated is returned to the main database.

In the database processing device (computer, information processing device, or program and object which perform the processing) according to the present invention, when execution of a particular processing 5 such as reorganization of the main database is instructed as an operation instruction, execution of the database processing for the main database as an object to be operated is stopped and the main database and its duplicate database are set to the quiescent 10 mode. By modifying the status in the DB definition information, the database to be operated is switched from the main database to the duplicate database, after which the quiescent mode of the main database and the duplicate database is released.

15 Next, a predetermined processing such as reorganization is performed to the main database which has become non-object to be operated by the status modification and a database access request is distributed to the duplicate data base which has become 20 an object to be operated by the status modification so as to continue database processing. Update information indicating the content of update processing performed to the duplicate database is output to the update information file.

25 When the predetermined processing to the main database is complete, the content of the main database is updated according to the update information output, thereby reflecting the update information on the main

data base. During this time, a business processing is continuously executed in the duplicate database and the update processing amount requiring reflection increases.

However, since the update reflection processing

- 5 executes update processing based on the update information in parallel, it is possible to execute at a higher speed than update by business and can catch up with the business processing.

When the reflection of the update information

- 10 is complete, execution of the database processing to the duplicate database as an object to be operated is stopped and the main database on which the update information is reflected and the duplicate database as an object to be operated are set to the quiescent mode.
- 15 By modifying the status in the DB definition information, the database to be operated is switched from the duplicate database to the main database, after which the stilled mode of the main database and the duplicate database is released and the database
- 20 processing by the main database is resumed.

- As has been described above, in the database processing device according to the present invention, when a predetermined processing should be performed to the main database, the predetermined processing can be
- 25 executed while continuing the database processing.

Other objects, features and advantages of the invention will become apparent from the following description of the embodiments of the invention taken

in conjunction with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 shows an outline of a database processing device according to a present embodiment.

5 Fig. 2 shows an outline of mode transition reorganized during database processing execution by switching from the normal operation mode to the replica operation mode of the present embodiment.

10 Fig. 3 shows an outline of the mode transition for returning from the replica operation mode to the normal operation mode after reorganization completion according to the present embodiment.

15 Fig. 4 shows an example program execution management information, table index storage information, and DB definition information according to the present embodiment.

Fig. 5 is a flowchart showing a replica operation start processing procedure according to the present embodiment.

20 Fig. 6 is a flowchart showing a replica operation end processing procedure according to the present embodiment.

25 Fig. 7 is a flowchart showing a DB update reflection processing procedure of step 603 according to the present embodiment.

Fig. 8 is a flowchart showing a procedure of a reflection control processing of step 704 and a

reflection execution processing of the update processing execution process according to the present embodiment.

Fig. 9 is a flowchart showing a procedure of
5 a an access object distribution process and access process according to the present embodiment.

Fig. 10 shows an example of update information record format according to the present embodiment.

10 Fig. 11 shows an example the execution process management table format of the update reflection process according to the present embodiment.

DESCRIPTION OF THE EMBODIMENTS

Description will now be directed to a
15 database processing system according to an embodiment of the present invention.

Fig. 1 shows an outline of a database processing system according to the present embodiment. As shown in Fig. 1, a database processing device 101 according to the present embodiment includes a database management processing section 102, an operation instruction section 104, a database operation processing section 110, a database access control section 111, a data staticizer 113, data de-staticizer 114, a dynamic copying section 115, a dynamic registration section, a generation management section 117, a current access generation setting section 118,

an access object setting section 122, an access object distribution section 123, an update information output section 124, and an update content reflection section 125. These processing sections 102, 104, 110-125 may 5 be hardware, programs, or objects where the respective processes are executed.

The database management processing section 102 receives a database access request from a program (business program and an application program) and the 10 operation instruction section 104 and controls operation of each processing section so as to execute database search, update, and reorganization.

The operation instruction section 104 receives an execution instruction of a particular 15 processing such as a database reorganization from a system administrator and instructs the database management processing section 102 to execute the particular processing. The database operation processing section 110 controls operations of the 20 dynamic copying section 115, the dynamic registration section 116, the generation management section 117, and the current access generation setting section 118.

The database access control section 111 controls operations of the data staticizer 113, the 25 data de-staticizer 114, the access object setting section 122, the access object distribution section 123, the update information output section 124, and the update content reflection section 125. The database

access control section 111 performs switching of a main database and a duplicate database, executes a predetermined processing in the main database not to be operated, continues the database processing in the 5 duplicate database, and reflects the update information file 109 on the main database.

The data staticizer 113 is a processing section for performing a processing to bring the main database to be operated and the duplicate database 10 holding the copy of it into the quiescent mode and to bring the main database in which the update information is reflected and the duplicate database to be operated into the quiescent mode.

The data de-staticizer 114 releases the 15 quiescent mode of the main database and the duplicate database. The dynamic copying section 115 is a processing section for copying the content of the main database onto the duplicate database. When the main database and the duplicate database holding its copy 20 are defined for the data base, the dynamic registration section 116 registers the definition content into the definition information file 106.

The generation management section 117 is a processing section for managing the operation mode of 25 the main database represented by a generation number "0" and the duplicate database represented by a generation number other than "0". The current access generation setting section 118 is a processing section

for switching the database represented by "current" and to be operated from the main database to the duplicate database or from the duplicate database to the main database.

5 The current database is a database allowed to access from the aforementioned program.

 The access object setting section 122 receives a request of accessing the database from the business AP 103 (program providing a business service 10 or program and object accessing the database), the operation instruction section 104, or the update content reflection section 125 and sets the access object on the main database or the duplicate database according to the access request content.

15 The access object distribution section 123 distributes the access request to the database set by the access object setting section 122 so as to distribute the database access request from the business AP 103, the execution instruction such as 20 reorganization from the operation instruction section 104, and the update information reflection request from the update content reflection section 125 to the main database or the duplicate database.

 The update information output section 124 25 outputs to the update information file 109 the update information indicating the content of the update processing performed on the duplicate database as an operation object.

Here, the main database is a copy origin database of the duplicate database. The update information is history information including at least one of the addition, deletion, and updating of a record 5 to/from/of the database. The update content reflection section 125 reflects the update information output to the update information file 109 to the main database where the aforementioned predetermined processing has been executed, by the parallel processing of a 10 plurality of execution processes corresponding to the key value in the update information.

The program for functioning the database processing device 101 as the database management processing section 102, the operation instruction 15 section 104, the database operation processing section 110, the database access control section 111, the data staticizer 113, the data de-staticizer 114, the dynamic copying section 115, the dynamic registration section 116, the generation management section 117, the current 20 access generation setting section 118, the access object setting section 122, the access object distribution section 123, the update information output section 124, and the update content reflection section 125 is recorded in a recording medium such as a CD-ROM 25 and stored on a magnetic disc, after which the program is loaded into a memory and is executed. It should be noted that the recording medium for recording the aforementioned program may be a recording medium other

than the CD-ROM. Moreover, the program may be installed from the recording medium to an information processing device or the program may be used by accessing the recording medium via a network.

5 Moreover, the disc system 105 includes a multiple write control section 126 for copying the main database 107 to create the duplicate database 108 and reflecting (also referred to as mirroring process) or not reflecting the update content of the main database
10 107 on the duplicate database 108 according to the instruction from the magnetic disc device containing the definition information file 106, the main database 107, the duplicate database 108, and the update information file 109, and the dynamic copying section
15 115.

The duplicate database creation processing utilizes the processing such as MRCP and shadow image which are the volume copy function in the disc device.

Fig. 2 shows an outline of a mode
20 transmission for reorganization during the database processing execution by switching from the normal operation mode to the replica operation mode according to the present embodiment. As shown in Fig. 2, in the database processing device according to the present
25 embodiment, the access request by the business AP 103 is processed by the main database if in the normal operation mode, and the content of the main database is copied to the duplicate database by the multiple write

control section 126. This multiple write processing may also be realized by the database processing section 102.

Here, when an instruction of a particular processing execution such as main database reorganization is input from the operation instruction section 104, the data staticizer 113 sets a lock mode for preventing execution of access processing to the database and outputting the data in the DB cache 1000 to the main database 107 (quiescent mode). Moreover, the generation management section 117 and the current access generation setting section 118 modify the mode of the DB definition information 1104 and the status content, releases the staticize by the data de-staticizer 114, and starts replica operation for switching the object to be operated from the main database to the duplicate database, thereby reorganizing the main database. The duplicate database outputs update information indicating the content of the update processing performed to the duplicate database by an access request processing from the business AP 103. The aforementioned reorganization processing and the access request processing may be executed in parallel and concurrently. Thus, by utilizing the quiescent mode, it is possible to switch the object to be operated from the main database to the duplicate database.

Thus, in the database processing device 101

of the present embodiment, the reorganization processing of the main database and the continuation of the database processing by the duplicate database are performed in parallel. Accordingly, when the main 5 database should be subjected to a predetermined processing such as reorganization, the database processing can be executed even when a predetermined processing is being performed.

Fig. 3 shows an outline of the mode 10 transition from the replica operation mode to the normal operation mode after the completion of reorganization in the present embodiment. As shown in Fig. 3, in the database processing device 101 of the present embodiment, when the reorganization of the main 15 database is complete, the aforementioned update information output is reflected on the main database, after which the object to be operated is switched from the duplicate database to the main database, thereby completing the replica operation and resuming the 20 normal operation by the main database.

In the aforementioned processing for reflecting the update information on the main database, the update information which has been output at the moment of the reflection is reflected on the main 25 database, after which the main database and the duplicate database to be operated are made into a quiescent mode and the update information output after the reflection until the staticize is reflected on the

main database, thereby reflecting the content of all the update processing performed for the duplicate database on the main database. It is also possible to set the quiescent mode when the number of the update 5 information processes has reached the number of the remaining processes and reflect the update information on the main database. After this, the normal operation mode by the main database is set in. By such a processing, the processing history of the update 10 processing to the duplicate database is provided as update information and it is possible to clarify the range of reflection on the main database.

Thus, in the database processing device 101 of the present embodiment, the reflection of update 15 information on the main database and the continuation of the database processing by the duplicate database are executed in parallel and the staticize is performed after the reflection of the update information is complete. Accordingly, it is possible to reduce the 20 time for waiting the access processing lock release.

Fig. 4 shows an example of program execution management information, table index storage information, and DB definition information according to the present embodiment. As shown in Fig. 4, the table index 25 storage information 1101 of the present embodiment contains a table/index name indicating a table and index name processed by an access request from the business AP 103, the operation instruction section 104,

or the update content reflection section 125 and the domain name of the database containing the tables and indexes.

Moreover, the DB definition information 1104
5 contains the domain name indicating the name of the domain as a substance of the database, the original domain name indicating the name of the domain in the main database, the generation indicating whether the database is a main database or a duplicate database,
10 the status indicating whether the database is an object to be operated, the mode indicating the operation mode of the database, and the access path name indicating the name of the access path to the database.

The program execution management information 1301 contains the program name indicating the name of the program which has performed an access request, an update reflection flag indicating whether the program performs an update reflection request, an access allowing flag indicating whether access in the replica 20 operation mode is allowed, a generation setting flag indicating whether a generation number is specified, and the generation number specified.

In the database processing device 101 of the present embodiment, when an execution instruction of a 25 particular processing such as database reorganization is input by a system administrator, the operation instruction section 104 receives the execution instruction input and passes the name of the table to

be subjected to the particular processing, the generation identifier of the database containing the table, a replica operation mode access allowance request requesting the allowance to access the database

5 in the replica operation mode, and the operation instruction instructing the execution of the particular processing, to the database access control section 111.

Upon reception of an instruction to execute a particular processing from the operation instruction
10 section 104, the database access control section 111 starts the replica operation with the duplicate database to be operated for the database as an object.

Fig. 5 is a flowchart showing a processing procedure of the replica operation start processing
15 according to the present embodiment. In step 501, the database access control section 111 searches the table/index name of the table index storage information 1101 by using the table name received from the operation instruction section 104 and acquires a
20 corresponding domain name. After this, the database access control section 111 searches from the DB definition information 1104 a record whose origin domain name coincides with the acquired domain name and whose generation is "0" and reads the mode in the
25 record to judge whether the database to be processed is in the normal operation mode. If the judgement results in other than the normal operation mode, control is passed to step 502 and the processing is terminated as

an error. If the judgement results in the normal operation mode, control is passed to step 503.

In step 503, the data staticizer 113 stores in a memory the information indicating that the main database to be processed and its duplicate database are in the quiescent mode, so that execution of the database processing for the main database to be operated is stopped and the main database and its duplicate database are staticized.

10 In step 504, the generation management section 117 sets information indicating the replica operation mode in the item of the record mode of the DB definition information 1104 searched and sets the main database in the replica operation mode.

15 In step 505, the generation management section 117 searches from the DB definition information 1104 a record whose origin domain name coincides with the acquired domain name and whose generation is other than "0", sets information indicating the replica operation mode in the item of the record mode searched, and sets the duplicate database to a replica operation mode.

20 In step 506, the database access control section 111 instructs the multiple write control section 126 of the disc system 105 to release the multiple write mode which is a mode for copying the content of the main database into the duplicate database.

In step 507, the current access generation setting section 118 sets "sub" indicating other than the database to be operated, in the status item in the record of the main database searched by the DB definition information 1104, sets "current" indicating a database to be operated, in the status item in the record of the duplicate database, and switches the database to be operated from the main database to the duplicate database.

10 In step 508, the data de-staticizer 114 deletes the information indicating the quiescent mode stored in the memory and releases the quiescent mode of the main database and the duplicate database.

As has been described above, when the replica operation is started, the database access control section 111 subjects the main database not to be operated to a particular processing such as reorganization instructed from the operation instruction section 104 and in parallel with this, continues the database processing in the duplicate database. Moreover, when the duplicate database to be operated is subjected to update processing, the update information output section 124 outputs the update information indicating the content of the update processing to the update information file 109. When the particular processing of the main database is complete, the database access control section 111 terminates the replica operation and performs

processing to return to the normal operation mode.

Fig. 6 is a flowchart showing a processing procedure of the replica operation end processing according to the present embodiment. In step 601, the 5 database access control section 111 reads out a record from the first entry of the DB definition information 110-4 and control is passed to step 602.

In step 602, if the generation inn the record read out is "0" and the status is "sub" and the mode is 10 "replica", i.e., if the record is a record of the main database where the particular processing has been executed in the replica operation mode, control is passed to step 603.

In step 603, the update content reflection 15 section 125 reflects the update information output to the update information file in the main database where the predetermined processing has been executed, by executing in parallel a plurality of execution processes corresponding to the key value in the update 20 information.

In step 604, the database access control section checks whether processes of all the entries in the DB definition information 1104 are complete. Unless processes of all the entries are complete, the 25 record of the next entry is read out and control is returned to step 602. If processes of all the entries are complete, the replica operation end processing is terminated.

Fig. 7 is a flowchart showing a processing procedure of the DB update reflection processing of step 603 in the present embodiment. In step 701, the update content reflection section 125 generates the number of update execution processes specified in advance by the system definition or the like.

In step 702, a record is read out from the head of the update information file 109 and control is passed to step 703. In step 703, the read out result of the record is checked to determine whether update information is present in the update information file 109. If update information is present, control is passed to step 704. Otherwise, control is passed to step 705.

In step 704, the update execution process is controlled so as to reflect the update information in the main database by the update information of the region to be processed.

In step 705, the data staticizer 113 stores in a memory the information indicating that the main database and its duplicate database are in the quiescent mode, so as to stop execution of the database processing for the duplicate database to be operated and set the main database and its duplicate database to the quiescent mode.

In step 706, a record is read out from the update information file 109 and the read out result is checked to determine whether update information is

present in the update information file 109. If the update information is present, control is passed to step 707. In step 707, the update execution process is controlled to reflect the update content by the update 5 information of the region to be processed.

In step 708, a record is read out from the update information file 109 and the read out result is checked to determine whether all the update information processes have been completed. Unless all the update 10 information processes are complete, control is returned to step 707. If all the update information processes are complete, control is passed to step 709. In step 709, all the execution processes generated in step 701 are instructed to be terminated.

15 In step 710, the current access generation setting section 118 searches the DB definition information 1104, sets "current" indicating the database to be operated in the status item in the record of the main database, sets "sub" indicating the 20 database not to be operated in the status item in the record of the duplicate database, and switches the database to be operated from the duplicate database to the main database.

 In step 711, the generation management 25 section 117 sets information indicating the normal operation mode in the item of the record mode of the main database searched and sets the main database into the normal operation mode.

In step 712, the generation management section 117 sets information indicating the non-operation mode in the item of the record mode of the duplicate database searched and sets the duplicate 5 database into the non-operation mode.

In step 713, the data de-staticizer 114 deletes the information indicating the quiescent mode stored in the memory and releases the quiescent mode of the main database and the duplicate database.

10 Fig. 8 is a flowchart showing a processing procedure of the reflection control processing and reflection execution of the update execution process of step 704. In step 801, the update content reflection section 125 reads in the key value of the database 15 record which has been updated from the update information in the record which has been read out.

In step 802, the entry of the execution process management table is decided by the value obtained as follows. The key value read in is divided 20 by the number of update execution processes and the obtained remainder is added by "1". The value thus obtained decides the entry.

In step 803, the execution process mode of the decided entry is read out from the execution 25 process management table. If the mode is other than ready, control is passed to step 804 for waiting until the mode becomes ready.

In step 805, the update information in the

record which has been read out is set in the update information of the decided entry. In step 806, the mode of the decided entry is modified to busy. In step 807, reflection execution is instructed to the 5 execution process identified by the process number of the entry decided.

In step 808, the update execution process checks whether processing instruction has been received from the update content reflection section 1125. If 10 the process instruction has not been received, control is passed to step 809 for waiting for the instruction.

In step 810, the instruction content received from the update content reflection section 125 is checked. If the instruction is other than the 15 instruction of processing end, control is passed to step 811. If the instruction is for the processing end, control is passed to step 813.

In step 811, the execution process management table is referenced and the update reflection processing is executed according to the update information stored in the entry of the update processing execution process. In step 812, the mode in the entry is returned to ready. On the other hand, in step 813, the processing of the update execution process is 25 terminated.

When the update reflection processing is started by the update execution process as has been described above, the access processes are distributed

according to the access object.

Fig. 9 is a flowchart showing a processing procedure of the access object distribution and the access processing according to the present embodiment.

- 5 In step 904, the table name in the access request is used to search a table/index name of the table index storage information 1101 and a corresponding domain name is acquired. In step 901, when the access object setting section 122 receives a database access request
- 10 from the business AP 103, the operation instruction section 104, or the update content reflection section 125, the access object setting section 122 checks whether the information indicating that the database to be accessed by the access request is in a quiescent
- 15 mode is stored in a memory and checks whether the database is locked (in quiescent mode). If it is locked, control is passed to step 902.

In step 902, the update reflection flag corresponding to the name of the program which has made

- 20 the access request is read out from the program execution management information 1301 and checked whether it is an update reflection request (update reflection flag is ON). If the access request is not the update reflection request, control is passed to
- 25 step 903 for waiting for lock release.

Step 905 checks whether specification of the generation identifier is present in the access request. If no specification of the generation identifier is

present, control is passed to step 906. If specification of the generation identifier is present, control is passed to step 907.

- Step 906 searches the DB definition information 1104 and acquires an entry in which the acquired domain name coincides with the original domain name and the status is "current". Moreover, step 907 searches the DB definition information 1104 and acquires an entry in which the acquired domain name coincides with the original domain name and the specified generation identifier coincides with the generation.

In step 908, the access object distribution section 123 references the mode in the entry thus obtained. If the database mode is a normal operation mode, control is passed to step 909. If the database mode is a replica operation mode, control is passed to step 910. If the database mode is a non-operation mode, control is passed to step 913.

Step 909 reads out the access path name in the entry and executes the access processing to the database according to the access path.

Step 910 checks whether the status in the entry is "current". If "current", control is passed to step 909. Otherwise, control is passed to step 911.

Step 911 reads out the access allowing flag corresponding to the name of the program which has made the access request from the program execution

management information 1301 and judges whether the access request is one allowed to access in the replica operation mode (access allowing flag is ON). If the access is allowed in the replica operation mode,
5 control is passed to step 909. Otherwise control is passed to step 912.

In step 912, an error is caused disabling access. Moreover, in step 913 also, an error is caused disabling access.

10 Fig. 10 shows an example of record format of the update information according to the present embodiment. As shown in Fig. 10, the update information of the present embodiment contains the domain name of the duplicate database where the update
15 processing is performed, the table name updated, the key value for identifying the database record where update processing is performed, the processing type indicating the update processing type, and the update data used for the update processing.

20 Fig. 11 is an example of the execution process management table format of the update reflection processing according to the present embodiment. As shown in Fig. 11, the execution process management table of the present embodiment contains the process number as information for identifying the update execution process (program and object), the processing mode indicating the mode of the process, and the update information used for the reflection

processing by the process.

As has been described above, in the database processing device according to the present embodiment, the database processing is continued by switching the 5 operation object to a plurality of databases, a predetermined processing is executed to the main database, and the update processing performed to the duplicate database is reflected to the main database which has been subjected to the predetermined 10 processing, after which the operation object is returned to the main data base. Accordingly, when it is necessary to perform a predetermined processing to the main database, the predetermined processing can be performed while continuing the database processing.

15 According to the present invention, when it is necessary to perform a predetermined processing to a database, it is possible to execute the database processing and the predetermined processing in parallel and concurrently.

20 It should be further understood by those skilled in the art that although the foregoing description has been made on embodiments of the invention, the invention is not limited thereto and various changes and modifications may be made without 25 departing from the spirit of the invention and the scope of the appended claims.